arXiv:1609.02302v1 [cs.NA] 8 Sep 2016

Soft Recovery Through $\ell_{1,2}$ Minimization with Applications in Recovery of Simultaneously Sparse and Low-Rank Matrices*

Axel Flinth†

Institut für Mathematik
Technische Universität Berlin

September 9, 2016

Abstract

This article provides a new type of analysis of a compressed-sensing based technique for recovering column-sparse matrices, namely minimization of the $\ell_{1,2}$-norm. Rather than providing conditions on the measurement matrix which guarantees the solution of the program to be exactly equal to the ground truth signal (which already has been thoroughly investigated), it presents a condition which guarantees that the solution is approximately equal to the ground truth. *Soft recovery statements* of this kind are to the best knowledge of the author a novelty in Compressed Sensing. Apart from the theoretical analysis, we present two heuristic proposes how this property of the $\ell_{1,2}$-program can be utilized to design algorithms for recovery of matrices which are sparse and have low rank at the same time.

Keywords: Compressed Sensing, Low-Rank matrices, $\ell_{1,2}$-minimization, Column Sparsity, Convex Optimization.

MSC2010: 52A41, 90C25.

1 Introduction

During the course of the last decade, the concept of *Compressed Sensing* [4] has gained an enormous interest in the signal processing community. Put shortly, Compressed Sensing is the science on how one can solve underdetermined equations using structural assumptions on the solutions (ground truth signals). Originally, the problem of recovering sparse vectors $x_0 \in \mathbb{R}^n$ from linear measurements Ax_0 were considered, but the ideas have been used to extend the theories to many other settings. In this paper, we will stay in the linear measurement regime, but consider another type of signal, namely *column-sparse matrices*. A matrix $X \in \mathbb{R}^{k,n}$ is thereby said to be *s-column sparse* when only s of its columns are non-zero. This problem appears naturally in a certain instance of the so called *blind deconvolution problem*. For motivational purposes, let us begin by describing this in a bit greater detail.

1.1 Blind Deconvolution and Column-sparse Matrices

Fundamentally, the problem of blind deconvolution reads as follows: From observing the convolution $v = w * x$ of two signals w and x, reconstruct w and x. Without any structurial assumptions on w and x, we cannot expect to succeed – this is already clear from considering the dimensions of the problem.

*This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

†E-mail: flinth@math.tu-berlin.de

There is a standard way of transforming the blind deconvolution to a matrix recovery problem [1]. It goes as follows: First, take the Fourier transform of the equation $v = w * x$:

$$\hat{v} = \hat{w} \odot \hat{x},$$

where $\odot$ refers to pointwise multiplication, i.e. $(\hat{w}_i \odot \hat{x})_i = \hat{w}_i \hat{x}_i$. Now we make two assumptions: The vector $\hat{w}$ lies in a k-dimensional subspace of $\mathbb{R}^n$ with basis $(d_i)_i$, and $\hat{x}$ is sparse in some basis $(e_i)_i$ of $\mathbb{R}^n$. Both the subspace and the basis are known. In applications, this could for instance mean that w is known to be bandlimited, and that x is a result of a user transmitting a short linear combination of some pre-determined main modes. Setting $\hat{w} = \sum_j h_j d_j$, and $\hat{x} = \sum_\ell z_\ell e_\ell$ and expanding our equation $\hat{v}_i = \hat{w}_i \hat{x}_i$, we obtain:

$$\hat{v}_i = \left(\sum_{j=1}^{k} h_j d_j \right)_i \left(\sum_{\ell=1}^{n} z_\ell e_\ell \right)_i = \sum_{j,\ell} d_j(i) e_\ell(i) h_j z_\ell = \mathcal{A}(hz^*), \quad i = 1, \dots n$$

where we defined the linear map $\mathcal{A}$ through its action on an $(k \times n)$-matrix

$$\mathcal{A}(M) = \sum_{j,\ell} d_j(i) e_\ell(i) M_{j,\ell}.$$

Hence, after proper reformulation, the *non-linear* blind deconvolution problem of recovering a *pair of vectors* can be seen as the problem of recovering a *matrix* from *linear* measurements $y = \mathcal{A}(M)$. It is also evident that matrices we want to recover, i.e. matrices of the form $Z_0 = h_0 z_0^*$ with $z_0 \in \mathbb{R}^n$ sparse and $h_0 \in \mathbb{R}^k$, are column sparse. One should also note that $h_0 z_0^*$ is low-rank - in fact, all columns are scaled versions of the vector h_0. Although we will primarily focus on the column-sparse structure in this work, we will also discuss the low rank aspect of the problem.

One could also imagine that one observes the sum of several convolutions with unknown filters: $v = \sum_{\ell=1}^{r} w_\ell * x_\ell$. One then sometimes speak of a *blind deconvolution and demixing* problem. A similar transformation of this problem leads to the task of recovering a matrix $Z_0 = \sum_{\ell=1}^{r} h_\ell^0 (z_\ell^0)^*$ from linear measurements. Assuming that the vectors z_ℓ^0, $\ell = 1, \dots r$ are sparse, Z_0 again becomes column sparse. Also, it trivially has rank at most r.

1.2 $\ell_{1,2}$-Minimization and Soft Recovery

Now let us return to the general problem of recovering a column-sparse matrix $Z_0 \in \mathbb{R}^{k,n}$ from linear measurements $y = \mathcal{A}(Z_0)$. A canonical approach [6] to solving this problem is to use $\ell_{1,2}$-minimization, i.e. to solve the problem

$$\min \|Z\|_{1,2} \text{ subject to } \mathcal{A}(Z) = y. \tag{$\mathcal{P}_{1,2}$}$$

The $\ell_{1,2}$-norm of a matrix is thereby defined as the sum of the Euclidean norm of its columns. There is much known about this problem – see for instance [6, 15].

Performing small numerical experiments with $\mathcal{A} : \mathbb{R}^{k,n} \to \mathbb{R}^m$ Gaussian, one can observe that $\mathcal{P}_{1,2}$ recovers the ground truth signals approximately already for values for m for which the solution $\widehat{Z}$ of $\mathcal{P}_{1,2}$ is almost never equal to Z_0. In particular, one notices that the solution $\widehat{Z}$ of of $\mathcal{P}_{1,2}$ seems to have the following behaviour:

- The directions of the columns $\widehat{Z}(i)$ of $\widehat{Z}(i)$ are well-aligned with the directions of the columns of $Z_0(i)$ (for the indices corresponding to non-zero columns of Z_0).
- The energy of $\widehat{Z}$ is concentrated on the columns corresponding to the non-zero columns of Z_0. That is: If $Z_0(i) = 0$, $\|\widehat{Z}(i)\|$ will be small, and vice versa.
- Also, if Z_0 has rank r, the subspace spanned by the first r left singular vectors of $\widehat{Z}$ is close to $\operatorname{ran} Z_0$.

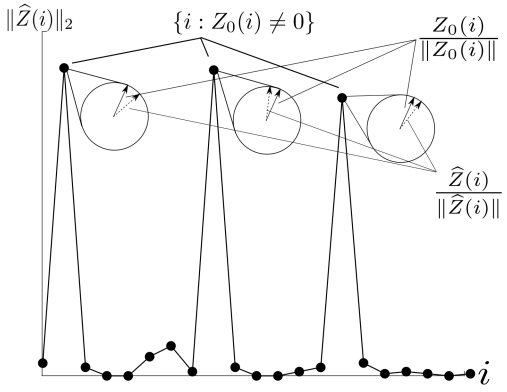

Figure 1: The aim of this work is to prove that the situation this figure depicts is the typical one. The vectors in the small circles show the directions of the columns of $Z_0(i)$ and $\widehat{Z}(i)$, respectively, for indices i with $Z_0(i) \neq 0$. Note that $k = 2$ in the figure.

See also Figure 1. We will say, in a consciously relatively unprecise manner, that Z_0 has been *softly recovered* by the program $\mathcal{P}_{1,2}$ if the solution $\widehat{Z}$ behaves as described above. The aim of this work is to provide a theoretical explanation of this feature of $\mathcal{P}_{1,2}$. In particular, we want to state conditions on $\mathcal{A}$ which imply soft recovery, but are weaker than the ones for exact recovery.

To the best knowledge of the author, this is the first time an analysis of this type has been conducted for any compressed sensing problem. One major reason for this is probably that the "original" Compressed Sensing problem of ℓ_1-minimization;

$$\min \|z\|_1 \text{ subject to } Az = y, \tag{$\mathcal{P}_1$}$$

does not exhibit soft recovery. In fact, Proposition A.1 in the Appendix of this paper shows for each matrix A and each ground truth signal z_0, either the solution of the problem $\mathcal{P}_1$ is equal to z_0, or there exists as at least one index i in the support of z_0 such that $\hat{z}(i) = 0$.

The reasons for conducting this analysis is twofold. First and foremost, we want to fundamentally understand the reason behind the soft recovery phenomenon. The technique used for the proof is also quite general, and could possibly be applied to other optimization-based recovery programs than $\mathcal{P}_{1,2}$ in the future.

The second reason is that if soft recovery is present, the solution $\widehat{Z}$ of $\mathcal{P}_{1,2}$ will give us information about the column-support and range of the ground truth solution Z_0 also when using too few measurements to recover it exactly. We will in fact use this fact to design two heuristic algorithms for recovery of column-sparse, low-rank matrices, and test their performance numerically.

The reader should not confuse soft recovery statements with statements about *instance optimality* [17]. The latter are statements of the form "if the ground truth signal is *approximately* column-sparse, $\mathcal{P}_{1,2}$ will *approximately* recover it". In contrast, we will consider *exactly* column-sparse ground truth signals, and investigate when they are *approximately* recovered.

1.3 Related Work on Low-Rank, Column-Sparse Matrix Recovery.

One should note that there are more sophisticated approaches to recovering low-rank and column-sparse matrices than $\ell_{1,2}$-minimization. For the sake of completeness, let us briefly describe two of these.

When recovering Z_0, we should try to take both the column sparsity and the low rank assumption into account. Since it is well known that low rank is promoted by the nuclear norm [14], defined as the sum of the singular

values of the matrix, a natural approach would be to minimize a weighted sum of the two norms

$$\min \|Z\|_{1,2} + \lambda \|Z\|_* \text{ subject to } \mathcal{A}(Z) = y.$$

This approach does perform reasonably well, but it has its flaws. From a practical point of view, choosing the correct λ is a non-trivial task. Also, the nuclear-norm minimization procedure tend to be a lot slower than procedures for minimizing norms like ℓ_1 or $\ell_{1,2}$. From a more theoretical standpoint, any such mixed-norm approach will need as least the same order of measurements to recover a low-rank, sparse matrix Z_0, as the minimum of the amounts needed to recover Z_0 only with help of $\|\cdot\|_{1,2}$ and $\|\cdot\|_*$, respectively. This 'single structure bottleneck' holds in much more generality - see [13]. The latter fact tells us that in order to develop a strategy which comes close to the optimal order $r(s+k)$(also see [13]), one has to deviate to other recovery methods. The authors of [13] provide a few examples of minimization problems which recover solutions from an, up to log-factors, optimal amount of measurements, but they are not computationally feasible.

The 'single-structure bottleneck' motivates why we in this paper choose to restrict the analysis to the program $\mathcal{P}_{1,2}$ - the inclusion of the nuclear norm term does not fundamentally change the performance, while it does make the analysis of the algorithm significantly more complicated.

An approach, which was proposed in [11], is to use an adapted version of the method of *power factorization*, introduced in [10]. Power factorization is fundamentally different from $\ell_{1,2}$-minimization, as it does not rely on minimizing one single convex function. The main idea is to exploit the fact that any rank-1-matrix can be written as $h_0 z_0^*$ explicitely: After initializing $\hat{h}$ and $\hat{z}$, one iteratively fixes one of the vector and then solves a least-square problem for the second:

$$\hat{h}^k = \text{argmin} \|\mathcal{A}(\hat{h}(\hat{z}^{k-1})^*) - b\|_2$$
$$\hat{z}^k = \text{argmin} \|\mathcal{A}(\hat{h}^{k-1}\hat{z}^*) - b\|_2.$$

In the column-sparse case, one incorporates the sparsity assumption when solving for $\hat{z}^k$. The method the authors of [11] choose is hard thresholding, but they state that other methods for sparse recovery could also be used for this purpose, for instance $CoSaMP$. They are able to prove that if one very carefully initializes the two iterates $\hat{h}$ and $\hat{z}$, the method converges with high probability to the correct solution when using $\gtrsim (s+k)\log(\max(n/s, k))$ measurements, which is essentially optimal. The initialization procedure used is however not computationally feasible. Because of this reason, the paper also provides an analysis for a more realistic initialization technique. The latter works under the same assumptions provided the ∞-norms of the vectors h_0 and z_0 are not too small (hence, they cannot be "well spread" on their respective supports). For general signals, they prove that $\gtrsim sk\log(n)$ measurements suffice.

During the final preparations of this paper, an updated version [11] was published, where an algorithm for matrices with rank higher than 1 was discussed. The idea is again to utilize the singular value decomposition $Z = U\Lambda V^*$ and solve alternating minimization problems. The authors prove that under some technical assumptions, any matrix Z_0 with singular value decomposition $U_0\Lambda_0 V_0^*$ with $U \in \mathbb{R}^{k,r}$ and $V \in \mathbb{R}^{r,n}$, where U is row-s_1-sparse and V is row-s_2-sparse satisfying some additional technical conditions (analogous to the $\|\cdot\|_\infty$-bounds from above) can with high probability be recovered with $Cr(s_1+s_2)\log(\max(ek/s_1, en/s_2)$ Gaussian measurements.

1.4 Outline

The remainder of this paper is organized as follows: In Section 2, we present notation, clarifies the context and and revise duality theory of convex optimization. The latter plays a crucial role in Section 3, where we present and prove the main findings of this work: We will provide results concerning all three aspects of soft recovery listed above. In particular, we will present a condition on $\mathcal{A}$ securing that angular distances between the columns of the solution of $\mathcal{P}_{1,2}$ and the ground truth signals are small, and provide upper bounds on how many Gaussian measurements are needed to secure that condition with high probability. In the final Section 4, we present two heuristic algorithms (called $NAST$ and Column Streamlining) for recovery of column-sparse, low rank matrices, and test their performance numerically.

2 Preliminaries

Let us begin by making the brief problem description from above a bit more precise, and introduce the notation which will be used throughout the paper. We consider a matrix $Z_0 \in \mathbb{R}^{n,k}$ (the ground truth signal) which is *column-sparse*, i.e. the set

$$\operatorname{supp} Z_0 = \{i \ \in [n] : Z_0(i) \neq 0\}$$

is small. $[n]$ is a shorthand for the set $\{1, 2, \dots n\}$ and $Z_0(i)$ denotes the i:th column of Z_0. The task is to recover the matrix from m linear measurements, given by the map $\mathcal{A} : \mathbb{R}^{n,k} \to \mathbb{R}^m$.

We will most often view the matrix Z_0 as a collection of n vectors $(Z_0(i))_{i\in[n]}$ in $\mathbb{R}^k$. Since we will make claims about the direction of these vectors, it will be convenient to decompose each of those vectors $Z(i)$ into a direction $h_i \in \mathbb{S}^{k-1}$ and a magnitude $z_i \in \mathbb{R}^+$. With this decomposition in mind, we write for $z \in \mathbb{R}^n_+$ and $H \in (\mathbb{S}^{k-1})^n$

$$z.H = \sum_{i\in[n]} z_i h_i e_i^* = \begin{bmatrix} z_1 h_1 & z_2 h_2 & \dots & z_n h_n \end{bmatrix}.$$

Note that if an element z_k of z is equal to zero, $z.H$ is well-defined even if the corresponding h_k is not specified. Also note that due to the fact that we have agreed to the convention hat $z \geq 0$, the composition $Z \to z.H$ is unique up to the fact that a column h_i can be choosen arbitrarily when $z_i = 0$.

Dual to the decomposition $Z \mapsto z.H$, we decompose the map $\mathcal{A}$ into the n maps $A_i : \mathbb{R}^k \to \mathbb{R}^m$ through

$$A_i h = \mathcal{A}(he_i^*).$$

Note that then $\mathcal{A}(z.H) = \sum_{i\in[n]} z_i A_i h_i$. We will sometimes also need the following induced map

$$A_H : \mathbb{R}^n \to \mathbb{R}^m, z \mapsto \mathcal{A}(z.H).$$

Note that the matrix representation of A_H is given by the matrix whose i:th column is given by $A_i h_i$. It is also not hard to convince oneself that the dual operator A_H^* is given by

$$A_H^* : \mathbb{R}^m \to \mathbb{R}^n, p \mapsto (\langle h_i, A_i^* p\rangle)_{i\in[n]}$$

We have already defined the $\ell_{1,2}$-norm. We will also need the $\ell_{\infty,2}$-norm:

$$\|Z\|_{\infty,2} = \max_{i\in[n]} \|Z(i)\|_2$$

Note that $\|\cdot\|_{1,2}$ harmonizes particularly well with our decomposition $Z \to z.H$; we have $\|Z\|_{1,2} = \|z\|_1 = \sum_{i\in[n]} z_i$ (this is true also for the $\ell_{\infty,2}$-norm, we will however never use that statement).

We will canonically measure distances between vectors $h, \hat{h} \in \mathbb{S}^{k-1}$ by their angular (geodesic) distance

$$\omega(h, \hat{h}) = \arccos(\langle h, \hat{h}\rangle).$$

ω defines a metric on $\mathbb{S}^{k-1}$. In particular, the triangular inequality holds:

$$\omega(h_1, h_2) \leq \omega(h_1, h_3) + \omega(h_3, h_2).$$

With this notation, we can formulate our main task as follows: Provide as weak conditions as possible guaranteeing that: ($Z_0 = z^0.H^0$ is the ground truth signal, $\widehat{Z} = \hat{z}.\widehat{H}$ is a solution of $\mathcal{P}_{1,2}$)

- $\omega(h_i^0, \hat{h}_i) \leq \alpha$ for $i \in S = \operatorname{supp} Z_0$ (for some previously specified, small α).
- $\|\hat{z}|_{S^c}\|_1 \leq \epsilon$, (for some previously specified, small ϵ).

2.1 Duality in Convex Programming and Exact Recovery.

Before considering the task of tackling the problem of soft recovery, let us deploy some standard techniques for formulating a condition guaranteeing that the ground truth signal Z_0 is the unique solution of $\mathcal{P}_{1,2}$. This will provide us with important insights of the structure of the $\ell_{1,2}$-norm, give us the opportunity to review basic facts about convex optimization, as well provide us with a benchmark that our soft recovery condition has to beat.

In the analysis, we will use *duality*. Given a convex program of the form

$$\min f(x) \text{ subject to } Ax = b,\ Dx \geq d \tag{$\mathcal{P}$}$$

where $f : \mathbb{R}^q \to \mathbb{R}$, $A \in \mathbb{R}^{\ell,q}$ $D \in \mathbb{R}^{\kappa,q}$, it is defined as follows: First, consider the Lagrangian

$$\mathcal{L} : \mathbb{R}^q \times \mathbb{R}^\ell \times \mathbb{R}^\kappa_+ \to \mathbb{R}, (x, \lambda, \mu) \mapsto f(x) + \langle \lambda, b - Ax\rangle + \langle \mu, d - Dx\rangle .$$

With the help of the Lagrangian, the dual function $g : \mathbb{R}^\ell \times \mathbb{R}^\kappa_+ \to \mathbb{R} \cup \{-\infty\}$ is defined through

$$g(\lambda, \mu) := \inf_{x \in \mathbb{R}^q} \mathcal{L}(x, \lambda, \mu).$$

The dual problem is then defined as

$$\max_{\lambda,\mu \geq 0} g(\lambda, \mu). \tag{$\mathcal{D}$}$$

The relation between the primal and dual problem is as follows: In any case (and in much greater generality than was presented here), the optimal value d^* of the dual problem is not greater than the optimal value of the primal problem p^*

$$p^* \geq d^*.$$

This is known as *weak duality*. Under quite general conditions, for instance for problems where f is convex, the constraints are linear , and that there exists an x with $b = Ax$ and $Dx \geq d$, we even have *strong duality*

$$p^* = d^*.$$

For more information about duality in convex optimization, we refer to the book [3].

Let us now state and prove the exact recovery condition.

Proposition 2.1. *Let $Z_0 = z^0.H^0$ be supported on the set S and $\mathcal{A} : \mathbb{R}^{n,k} \to \mathbb{R}^m$ be given. Z_0 is a solution to the program $\mathcal{P}_{1,2}$ if and only if there exists a $p \in \mathbb{R}^m$ such that $V = \mathcal{A}^* p$ satisfies*

$$V(i) = h^0_i, \quad i \in S \tag{1}$$

$$\|V(i)\|_2 \leq 1, \quad i \notin S \tag{2}$$

Proof. Let us begin by calculating the dual problem $\mathcal{D}_{1,2}$. The Lagrange function is given by

$$\mathcal{L} : \mathbb{R}^{n,k} \times \mathbb{R}^m \to \mathbb{R}, (Z, p) \mapsto \|Z\|_{1,2} + \langle p, b - \mathcal{A}(Z)\rangle = \langle p, b\rangle + \sum_{i \in n} z^0_i (1 - \left\langle A^*_i p, h^0_i \right\rangle).$$

From this formula, it is easily seen that the dual function g is given by

$$g(p) = \begin{cases} \langle p, b\rangle \text{ if } \|(A^*_i p)_{i\in[n]}\|_{\infty,2} \leq 1 \\ -\infty \text{ else.} \end{cases}$$

Due to strong duality, Z_0 is a solution to $\mathcal{P}_{1,2}$ if and only if the optimal value of the problem

$$\max \langle p, b\rangle \text{ subject to } \forall i : \|A^*_i p\|_2 \leq 1$$

is given by $\|Z_0\|_{1,2} = \|z^0\|_1$. Due to the compactness of the set of feasible vectors, this is the case if and only if there exists a $p \in \mathbb{R}^m$ with $\|A_i^* p\|_2 \leq 1$ and $\langle p, b\rangle = \|z^0\|_1$. It is however clear that in this case,

$$\|z^0\|_1 = \langle p, b\rangle = \sum_{i\in S} z_i^0 \left\langle p, A_i h_i^0 \right\rangle = \sum_{i\in S} z_i^0 \left\langle A_i^* p, h_i^0 \right\rangle$$

if and only if $A_i^* p = h_i^0$. Noting that $(\mathcal{A}^* p)(i) = A_i^* p$, we see that the claim has been proven. □

We will tackle also the soft recovery problem using the tools of duality, however not as directly as above. The main idea will be to view the $\ell_{1,2}$-minimization as a family of ℓ_1-minimizations . To be precise, let us assume that the solution of $\mathcal{P}_{1,2}$ is given by $\hat{z}.\widehat{H}$. It is then not hard to convince oneself that $\hat{z}$ is the solution of the problem

$$\|z\|_1 = \sum_{i\in[n]} z_i \text{ subject to } A_{\widehat{H}} z = b, z \geq 0 \qquad (\mathcal{P}_1^+(\widehat{H}))$$

After all, for any other $z \geq 0$ with $A_{\widehat{H}} z = b$, we have $\mathcal{A}(z.\widehat{H}) = b$ and, due to the optimality of $\hat{z}.\widehat{H}$, $\|\hat{z}\|_1 = \|\hat{z}.\widehat{H}\|_{1,2} \leq \|z.\widehat{H}\|_{1,2} = \|z\|_1$. Using this relatively elementary observation, we can prove the following property of the minimizer $\hat{z}.\widehat{H}$

Lemma 2.2. *Let $\hat{z}.\widehat{H}$ be the minimizer of $\mathcal{P}_{1,2}$ for $b = \mathcal{A}(z^0.H^0)$, where $\|z^0\|_1 = 1$. Then*

$$\min_{\langle b,p\rangle \geq 1} \max_{i\in[n]} \langle \hat{h}_i, A_i^* p\rangle \geq 1$$

Proof. The dual problem of $\mathcal{P}_1^+(\widehat{H})$ can be written in the following way

$$\max \langle p, b\rangle \text{ subject to } \max_{i\in[n]} (A_{\widehat{H}}^* p)_i \leq 1. \qquad (\mathcal{D}_1^+(\widehat{H}))$$

(The, relatively standard, calculations leading up to this formulation of the dual problem are, for completeness, presented in Lemma A.2 in the Appendix.) There exists a $\hat{z} \geq 0$ so that $\hat{z}.\widehat{H}$ solves $\mathcal{P}_{1,2}$, i.e. in particular $b = \mathcal{A}(\hat{z}.\widehat{H}) = A_{\widehat{H}}\hat{z} = b$. Hence, strong duality holds. Therefore, the optimal value of $\mathcal{D}_1(\widehat{H})$ is equal to $\|\hat{z}\|_1$.

Now, due to the optimality of $\hat{z}.\widehat{H}$,

$$\|\hat{z}\|_1 = \|\hat{z}.\widehat{H}\|_{1,2} \leq \|z^0.H^0\|_{1,2} = \|z^0\|_1 = 1.$$

This means that the optimal value of $\mathcal{D}_1(\widehat{H})$ is not larger than one, which implies that there cannot exist a p with $\max_{i\in[n]} \langle h_i, A_i^* p\rangle = \max_{i\in n}(A_{\widehat{H}}^* p)_i < 1$ and $\langle p, b\rangle \geq 1$. If that would be the case, a renormalized version of p would satisfy $\max_{i\in[n]}(A_{\widehat{H}}^* p)_i \leq 1$ and $\langle b, p\rangle > 1$, which would contradict the fact that the optimal value of $\mathcal{D}_1(\widehat{H})$ is not larger than 1. Hence, the claim has been proven. □

3 Main Results

With Lemma 2.2 in our toolbox, we are ready to state and prove one of the main result of this paper. It provides a condition guaranteeing that the angular distances between the columns of the solution of $\mathcal{P}_{1,2}$ and the ground truth signals are small.

Theorem 3.1. *Let $Z_0 = z_0.H_0$ be supported on the set S , $i^* \in S$ and $\alpha > 0$ be fixed. If there exists a vector $p^* \in \mathbb{R}^m$ so that $V = \mathcal{A}^* p$ has the following properties*

$$\sum_{i\in S} \left\langle h_i^0, V(i)\right\rangle z_i^0 \geq \|z^0\|_1 \qquad (3)$$

$$\|V(i)\|_2 < 1 \quad i \neq i^* \qquad (4)$$

$$\omega\left(h_{i^*}^0, \tfrac{V(i^*)}{\|V(i^*)\|_2}\right) \leq \alpha. \qquad (5)$$

$$\|V(i^*)\|_2 \cos\alpha \leq 1 \qquad (6)$$

Then any minimizer $\hat{z}.\widehat{H}$ of $\mathcal{P}_{1,2}$ obeys

$$\omega(\hat{h}_{i^*}, h^0_{i^*}) \leq 2\alpha.$$

Proof. Let us begin by noting that we without loss of generality may assume that $\|z^0\|_1 = 1$. This since only the norms of the columns of the minimizer changes when we scale z^0, and not the directions of them.

We have, due to (3)

$$\langle b, p^* \rangle = \langle \mathcal{A}(z^0.H^0), p^* \rangle = \sum_{i \in S} z^0_i \left\langle A_i h^0_i, p^* \right\rangle \geq 1.$$

This, together with Lemma 2.2 implies that there exists an $i \in [n]$ with

$$\langle A^*_i p^*, \hat{h}_i \rangle \geq 1.$$

Due to (4), this index i must be equal to i^*.

Now suppose that $\omega(\hat{h}_{i^*}, h^0_{i^*})$ is strictly larger than 2α. Then by (5) and the triangle inequality of ω

$$\omega(\hat{h}_{i^*}, \tfrac{A^*_i p^*}{\|A^*_i p^*\|}) \geq \omega(\hat{h}_{i^*}, h^0_{i^*}) - \omega\left(h^0_{i^*}, \tfrac{A_{i^*} p^*}{\|A_{i^*} p^*\|_2}\right) > 2\alpha - \alpha = \alpha,$$

which implies

$$\langle A^*_{i^*} p^*, \hat{h}_{i^*} \rangle = \|A^*_{i^*} p^*\|_2 \cos(\omega(\hat{h}_{i^*}, \tfrac{A^*_{i^*} p^*}{\|A^*_{i^*} p^*\|})) \leq \|A^*_{i^*} p^*\|_2 \cos(\alpha) \leq 1$$

due to (6). This is a contradiction and the proof is finished. □

Let us make some comments on (3)-(6).

First, the set defined by the four constraints is convex for $\alpha \leq \frac{\pi}{2}$. First, it is easily seen that both (3), (4) and (6) define convex sets. As for (5), notice that it can be rewritten as

$$\langle A^*_{i^*} p, h^0_{i^*} \rangle \geq \cos(\alpha) \|A^*_i p\|_2$$

If $\alpha \leq \frac{\pi}{2}$, $\cos(\alpha) \geq 0$ and hence for p_1, p_2 satisfying (5) and $\theta \in (0,1)$, we have

$$\begin{aligned} \|A^*_{i^*}(\theta p_1 + (1-\theta)p_2)\|_2 \cos\alpha &\leq \theta \|A^*_{i^*} p_1\|_2 \cos\alpha + (1-\theta)\|A_{i^*} p_2\|_2 \cos\alpha \\ &\leq \theta \left\langle A_{i^*} p_1, h^0_{i^*} \right\rangle + (1-\theta) \left\langle A_{i^*} p_2, h^0_{i^*} \right\rangle = \left\langle A^*_{i^*}(\theta p_1 + (1-\theta)p_2, h^0_{i^*} \right\rangle. \end{aligned}$$

Second, the condition described in Theorem 3.1 is in essence weaker than the one described in Proposition 2.1, at least when dealing with reasonable random matrices. To see this, let us first note that for such matrices, it does not matter if we replace the strict inequality sign in (4) with a "$\leq$"- this does not change for instance the Gaussian width [9] or the statistical dimension [2] of the set (see also below). With this slight change, we see that a vector satisfying (1) and (2) necessarily obeys

$$\begin{aligned} \sum_{i \in S} \left\langle h^0_i, A^*_i p^* \right\rangle z^0_i = \sum_{i \in S} \left\langle h^0_i, h^0_i \right\rangle z^0_i &= \|z^0\|_1 \\ \|A^*_i p\|_2 \leq 1, \quad i \neq i^* \\ \omega(h^0_{i^*}, \tfrac{A^*_{i^*} p^*}{\|A^*_{i^*} p^*\|}) = \omega(h^0_{i^*}, h^0_{i^*}) &= 0 \leq \alpha \\ \|A_{i^*} p\|_2 \cos(\alpha) = \|h^0_{i^*}\|_2 \cos(\alpha) &= \cos(\alpha) \leq 1. \end{aligned}$$

Hence, if the exact recovery statement is fulfilled, the soft recovery statement (in essence) is also.

Third, it is not a big problem that we for every $i \in S$ need to (separately) ensure the existence of one vector in the range of $\mathcal{A}^*$ fulfilling 3-(6) to ensure that $\omega(h^0_i, \hat{h}^0_i)$ for all $i \in S$. To argue why this is the case, consider this

(slightly reformulated) inequality from [2, Theorem 7.1]: for a convex cone $C \subseteq \mathbb{R}^d$, we have for an m-dimensional uniformly distributed random subspace V

$$m \geq d - \delta(C) + \sqrt{4d \log \eta^{-1}} \Rightarrow \mathbb{P}\left(C \cap V \neq 0)\right) \leq \eta, \tag{7}$$

where $\delta(C)$ denotes the *statistical dimension* of the cone C. Due to the linear structure of $\operatorname{ran} \mathcal{A}^*$, there exists a vector in $\operatorname{ran} \mathcal{A}^*$ fulfilling 3-6 exactly when $\operatorname{ran} \mathcal{A}^* \cap \mathcal{C}_{i^*} \neq \{0\}$, where $\mathcal{C}_i$ is the cone generated by the constraints (3)-(6). (7) now tells us that if V is a uniformly distributed random subspace of dimension m in $\mathbb{R}^{k,n}$

$$\begin{aligned}\mathbb{P}\left(\forall i \in S : \mathcal{C}_i \cap V \neq 0)\right) &= 1 - \mathbb{P}\left(\exists i \in S : \mathcal{C}_i \cap V = 0)\right) \geq 1 - \sum_{i \in S} \mathbb{P}\left(\mathcal{C}_i \cap V = 0\right) \\ &\geq 1 - s \max_{i \in S} \mathbb{P}\left(\mathcal{C}_i \cap V = 0\right) \geq 1 - \eta\end{aligned}$$

if $m \geq nk - \min_i \delta(\mathcal{C}_i) + \sqrt{4d \log((s\eta)^{-1})}$. In the case that $\mathcal{A}$ is a random Gaussian, $\operatorname{ran} \mathcal{A}$ will have exactly the distribution of V above. Hence, the number of Gaussian measurements needed to secure the existence of a p satisfying (3)-(6) for all $i \in S$ is basically the same amount as the one needed to secure one satisfying the "hardest" constraint. The only difference is a factor of log-type.

3.1 The Statistical Dimension of the Cone Generated by (3) - (6).

We have already discussed the concept of *statistical dimension* $\delta(C)$ of a convex cone $C \subseteq \mathbb{R}^d$. It was introduced in [2]. The importance of the concept is captured in the formula (7) - the statistical dimension provides a threshold amount of Gaussian measurements needed for $\operatorname{ran} A$ to intersect the cone. It can be defined in many ways. The most convenient way is probably to first define it for closed convex cones C as

$$\delta(C) = \mathbb{E}\left(\|\Pi_C g\|_2^2\right),$$

where Π_C denotes the (non-linear) orthogonal projection onto C, and $g \in \mathbb{R}^d$ a Gaussian vector. For general convex cones K, we define $\delta(K)$ as the statistical dimension of the closure of K.

It is hard to give an exact expression for the statistical dimension of the cone generated by (3) - (6) (let us denote it by $\mathcal{C}_{i^*}$, as we did above). In this section, we will prove a lower bound of $\delta(\mathcal{C}_{i^*})$, which through (7) provides an upper bound on the amount of measurements needed to guarantee soft recovery. This bound is also not a closed expression, but one that easily can be calculated numerically. The first step of the argumentation will be to identify a subset of $\mathcal{C}_{i^*}$ which is a convex cone whose statistical dimension is easier to handle. Let us begin by describing the subset.

Lemma 3.2. *Let $\alpha < \frac{\pi}{2}$. Then, $\mathcal{C}_{i^*}$ contains $\operatorname{cone} \mathcal{M}_{i^*}$, where $\mathcal{M}_{i^*} \subseteq \mathbb{R}^{k,n}$ is given by the equations*

$$\begin{aligned}\|V_i\|_2 &< 1 \quad i \neq i^* \\ \left\langle V_i, h_i^0 \right\rangle &\geq \sigma \quad i \in S \backslash \{i^*\} \\ \|V_{i^*}\|_2 &\leq \tfrac{1}{\cos\alpha} \\ \left\langle V_{i^*}, h_{i^*}^0 \right\rangle &\geq \tfrac{\sigma}{\cos\alpha}.\end{aligned}$$

σ is a parameter defined through

$$\sigma := \frac{1}{1 + (\frac{1}{\cos(\alpha)} - 1)\frac{z_i^0}{\|z^0\|_1}}.$$

Proof. Let $V \in \mathcal{M}_{i^*}$. Then trivally (4) and (6) are satisfied. As for (3), we have

$$\sum_{i \in S} \left\langle h_i^0, V_i \right\rangle z_i^0 \geq z_{i^*}^0 \cdot \tfrac{\sigma}{\cos\alpha} + \sum_{i \in S \backslash i^*} \sigma z_{i^*}^0 = z_{i^*}^0 \cdot \tfrac{\sigma}{\cos\alpha} + \sigma(\|z^0\|_1 - z_{i^*}^0) = \|z^0\|_1.$$

As for (5), we need to prove that $\left\langle V_{i^*}, h_{i^*}^0 \right\rangle \geq \|V_{i^*}\|_2 \cos\alpha$. For this, it suffices to show that $\sigma \geq \cos\alpha$, since we already know that $\left\langle V_{i^*}, h_{i^*}^0 \right\rangle \geq \frac{\sigma}{\cos\alpha}$ and $\|V_{i^*}\|_2 \leq \frac{1}{\cos\alpha}$. This is however not hard: $\sigma \geq \cos\alpha$ is equivalent to

$$1 \geq \cos\alpha + \tfrac{z_i^0}{\|z^0\|_1}(1-\cos\alpha) = \tfrac{z_i^0}{\|z^0\|_1} + \cos\alpha\left(1 - \tfrac{z_i^0}{\|z^0\|_1}\right),$$

which is true, since $\cos\alpha \leq 1$.

□

With the help of the previous lemma and the definition of the statistical dimension of a cone, we derive the following lower bound for $\delta(\mathcal{C}_{i^*})$:

$$\delta(\mathcal{C}_{i^*}) \geq \delta(\operatorname{cone}\mathcal{M}_{i^*}) = \mathbb{E}\left(\|\Pi_{\operatorname{cone}\mathcal{M}_i} G\|_2^2\right) = nk - \mathbb{E}\left(\inf_{\tau>0}\operatorname{dist}(G-\tau\mathcal{M}_{i^*})^2\right) \geq nk - \inf_{\tau>0}\mathbb{E}\left(\operatorname{dist}(G-\tau\mathcal{M}_{i^*})^2\right), \tag{8}$$

where $G \in \mathbb{R}^{k,n}$ is Gaussian. We used the Lemma of Fatou and the following trick from [2]: for $g \in \mathbb{R}^d$ Gaussian and $C \subseteq \mathbb{R}^d$, we have

$$\mathbb{E}\left(\|\Pi_C g\|_2^2\right) = \mathbb{E}\left(\|g\|_2^2 - \min_{c\in C}\|g-c\|_2^2\right) = d - \mathbb{E}\left(\min_{c\in C}\|g-c\|_2^2\right).$$

In other words, the task of bounding $\delta(\mathcal{M}_{i^*})$ from below can be accomplished by bounding $\inf_{\tau>0}\mathbb{E}\left(\operatorname{dist}(G-\tau\mathcal{M}_{i^*})^2\right)$ above. In order to solve this task, let us first have a look at the function $V \to \operatorname{dist}(V-\tau\mathcal{M}_{i^*})^2$.

Lemma 3.3. *Define β through $\sigma = \cos(\beta)$ and let $V \in \mathbb{R}^{k,n}$ Then*

$$\operatorname{dist}(V-\tau\mathcal{M}_{i^*})^2 = \sum_{i\in[n]} d_i(V_i),$$

where $d_i : \mathbb{R}^n \to \mathbb{R}$ are defined through

$$d_i(v) = \operatorname{pos}(\|v\|_2 - \tau)^2 \qquad i \notin S$$

$$d_i(v) = \begin{cases} (\langle v, h_i^0\rangle - \tau\cos\beta)^2 + \operatorname{pos}(\|\Pi_{\langle h_i^0\rangle^\perp} v\|_2 - \tau\sin\beta)^2 & \text{if } \|\Pi_{\langle h_i^0\rangle^\perp} v\|_2 \geq \tan\beta\langle h_i^0, v\rangle \\ & \text{or } \left\langle v, h_i^0\right\rangle \leq \tau\cos\beta \\ \operatorname{pos}(\|v\|_2-\tau)^2 & \text{else.} \end{cases} \qquad i \in S\backslash\{i^*\}$$

$$d_{i^*}(v) = \begin{cases} (\langle v, h_i^0\rangle - \tau\frac{\cos\beta}{\cos\alpha})^2 + \operatorname{pos}(\|\Pi_{\langle h_i^0\rangle^\perp} v\|_2 - \tau\frac{\sin\beta}{\cos\alpha})^2 & \text{if } \|\Pi_{\langle h_i^0\rangle^\perp} v\|_2 \geq \tan\beta\langle h_i^0, v\rangle \\ & \text{or } \left\langle v, h_i^0\right\rangle \leq \tau\frac{\cos\beta}{\cos\alpha} \\ \operatorname{pos}(\|v\|_2 - \frac{\tau}{\cos\alpha})^2 & \text{else.} \end{cases}$$

Proof. First, it is immediately clear that

$$\operatorname{dist}(V-\tau\mathcal{M}_{i^*})^2 = \min_{\substack{\|w\|_2\leq\frac{\tau}{\cos\alpha} \\ \langle w,h_i^0\rangle\geq\tau\frac{\cos\beta}{\cos\alpha}}} \|v_{i*}-w\|_2^2 + \sum_{i\in S\backslash\{i^*\}} \min_{\substack{\|w\|_2\leq\tau \\ \langle w,h_i^0\rangle\geq\tau\cos\beta}} \|v_i-w\|_2^2 + \sum_{i\notin S}\min_{\|w\|_2\leq\tau}\|v_i-w\|_2^2.$$

It is also easy to see that $\min_{\|w\|_2\leq\tau}\|v_i-w\|_2^2 = \operatorname{pos}(\|v_i\|_2-\tau)^2$ (we can choose w parallel to v_i and exhaust as much of its length as possible - in particular its whole length if $\|v_i\|_2 \leq \tau$). For the other expressions, we have to be a little more careful. We will only treat the case that $i \in S\backslash\{i^*\}$ - the case $i = i^*$ is analogous (the only difference are extra $\frac{1}{\cos\alpha}$-factors appearing everywhere, whence we choose to omit that case).

To calculate

$$\min_{\substack{\|w\|_2\leq\tau \\ \langle w,h_i^0\rangle\geq\tau\cos\beta}} \|v-w\|_2^2,$$

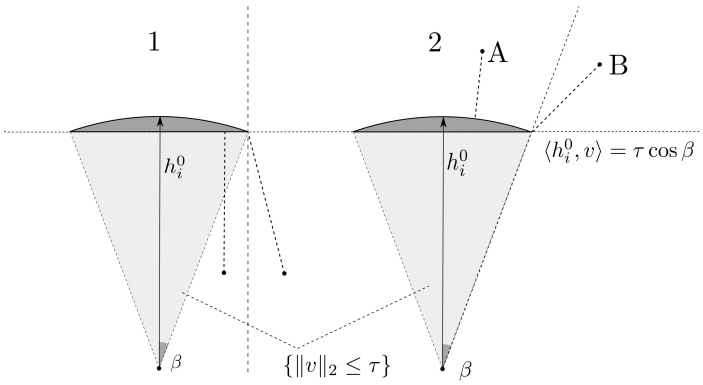

Figure 2: Graphical depiction of the argumentation in the Proof of Lemma 3.3

.

we distinguish two cases. A graphical depiction of the argumentation in each case can be found in Figure 2

Case 1 : $\left\langle v, h_i^0\right\rangle \leq \tau\cos(\beta)$**.** In this case, the minimizer of $\|v-w\|_2$ must lie on the $(k-1)$-dimensional disc

$$\left\{w \mid \left\langle w, h_i^0\right\rangle = \tau\cos\beta,\ \|\Pi_{\left\langle h_i^0\right\rangle^\perp} w\|_2 \leq \tau\sin\beta\right\}.$$

With this insight, it is clear that the minimum is given as claimed.

Case 2: $\left\langle v, h_i^0\right\rangle \geq \tau\cos(\beta)$**.** In this case, the minimizer must lie on the spherical cap $\|w\|_2 = 1, \omega(w, h_i^0) \leq \beta$. This makes it clear that the minimum depends on the angular distance between $\frac{v}{\|v\|_2}$ and h_i^0.

If the angle is smaller than β, we can choose w parallel to v, and the smallest possible separation is equal to $\mathrm{pos}(\|v\|_2 - \tau)$. In the other case, we may choose w in the plane spanned by h_i^0 and v with an angular distance to h_i^0 at most β. The minimizer is thus equal to $\tau(\cos\beta\ h_i^0 + \sin\beta\ \Pi_{\left\langle h_i^0\right\rangle^\perp} v)$, and the minimum equal to the claimed expression. □

From the proof of the last lemma, it is evident that we can estimate for $i \in S\setminus\{i^*\}$

$$d_i(v) \leq (\langle v, h_i^0\rangle - \tau\cos\beta)^2 + \mathrm{pos}(\|\Pi_{\left\langle h_i^0\right\rangle^\perp} v\|_2 - \tau\sin\beta)^2 \tag{9}$$

This may seem crude, but we in fact only make a non-strict estimate on the set $\left\{v|\omega(\frac{v}{\|v\|_2}, h_i^0) < \beta, \left\langle v, h_i^0\right\rangle \geq \tau\cos(\beta)\right\}$. This set has a very small Gaussian measure for small angles β and large dimensions k. An analogous estimate can be made for $i = i^*$.

This final estimate allows us to prove the following proposition.

Proposition 3.4. *Let $G \in \mathbb{R}^{k,n}$ be Gaussian. We have*

$$\mathbb{E}\left(\mathrm{dist}(G - \tau\mathcal{M}_i)^2\right) \leq \Phi_{k,s,n,\alpha,\beta}(\tau)$$

Where

$$\Phi_{k,s,n,\alpha,\beta}(\tau) := s + \tau^2\tfrac{\cos^2\beta}{\cos^2\alpha} + \mathbb{E}\left(\mathrm{pos}(\|g^{k-1}\|_2 - \tau\tfrac{\sin\beta}{\cos\alpha})^2\right) + (s-1)\left(\tau^2\cos^2\beta + \mathbb{E}\left(\mathrm{pos}(\|g^{k-1}\|_2 - \tau\sin\beta)^2\right)\right) + (n-s)\mathbb{E}\left(\mathrm{pos}(\|g^k\|_2 - \tau)^2\right)$$

where g^d denotes an d-dimensional Gaussian and $s = |S|$.

In particular, if $\mathcal{A} : \mathbb{R}^{k,n} \to \mathbb{R}^m$ is Gaussian with $m \geq \inf_{\tau>0}\Phi_{k,s,n,\alpha,\beta}(\tau)$, 2α-soft recovery of the i^*-column is guaranteed with high probability.

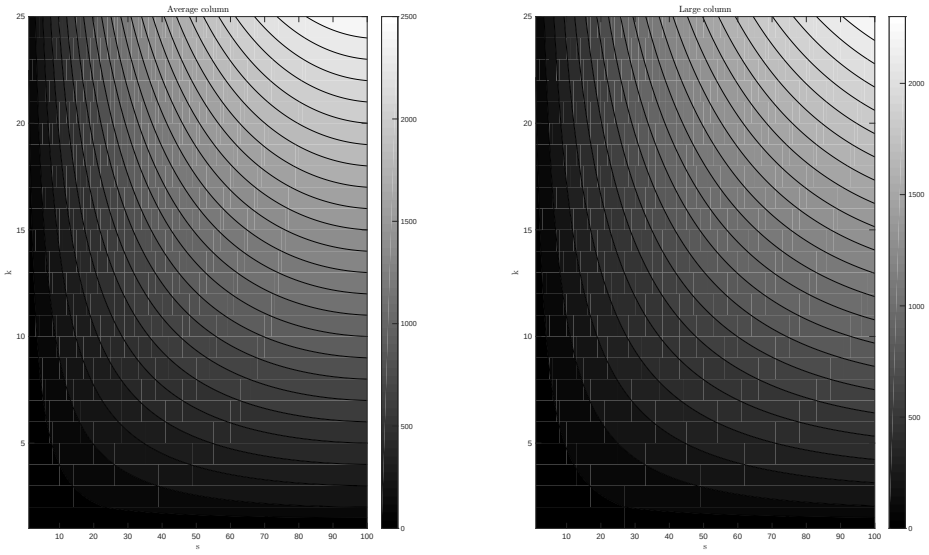

Figure 3: $m_{k,s,n,\alpha,\beta}$ depending on k and s for average (left) and large (right) column. $n = 100$, $\alpha = \frac{\pi}{10}$.

Proof. We have due to Lemma 3.3

$$\mathbb{E}\left(\operatorname{dist}(G - \tau\mathcal{M}_i)^2\right) = \sum_{i\in[n]} \mathbb{E}\left(d_i(G_i)\right) = \mathbb{E}\left(d_{i^*}(G_{i^*})\right) + \sum_{i\in S\setminus\{i^*\}} \mathbb{E}\left(d_i(G_i)\right) + (n-s)\mathbb{E}\left(\operatorname{pos}(\|g^k\|_2 - \tau)^2\right),$$

since each column of a Gaussian $G \in \mathbb{R}^{k,n}$ is a k-dimensional Gaussian. (9) furthermore implies for $i \in S\setminus\{i^*\}$

$$\begin{aligned}\mathbb{E}\left(d_i(G_i)\right) \leq \mathbb{E}\left(\langle G_i, h_i^0\rangle^2\right) - 2\tau\cos\beta\ \mathbb{E}\left(\langle G_i, h_i^0\rangle\right) + \tau^2\cos^2\beta + \mathbb{E}\left(\operatorname{pos}(\|\Pi_{\langle h_i^0\rangle^\perp} G_i\|_2 - \tau\sin\beta)^2\right)\\ = 1 + \tau^2\cos^2\beta + \mathbb{E}\left(\operatorname{pos}(\|g^{k-1}\|_2 - \tau\sin\beta)^2\right),\end{aligned}$$

where in the last step we used that projections of Gaussians again are Gaussians. By performing a similar calculation for $\mathbb{E}\left(d_{i^*}(G_{i^*})\right)$ and summing all of the terms, we arrive at the first statement.

The second statement is now a trivial consequence of the first one together with the lower bound (8) and the statement (7).

□

Remark 3.5. *By putting $\alpha = 0$ (which implies $\beta = 0$), we arrive at a corresponding formula for the amount of measurements needed for exact recovery. Going through all of the calculations in this section again, one sees that in the case $\alpha = 0$, all inequalities are in fact equalities, and we hence arrive at the exact value for the statistical dimension in this case. We may even simplify (since $\mathbb{E}\left(\operatorname{pos}(\|g^{k-1}\|_2 - \tau\cdot 0)^2\right) = k-1$)*

$$\Phi_{k,s,n,0,0}(\tau) = sk + s\tau^2 + (n-s)\mathbb{E}\left(\operatorname{pos}(\|g^k\|_2 - \tau)^2\right).$$

Using computer software, one can calculate $m_{k,s,n,\alpha,\beta} = \inf_{\tau>0}\Phi_{k,s,n,\alpha,\beta}(\tau)$. Let us discuss some of the qualitative behaviours one can read out of these numerical evaluations.

In Figure 3, we fix $\alpha = \frac{\pi}{10}$, $n = 100$ and plot $m_{k,s,n,\alpha,\beta}$ for varying k and s. We investigate two situations: In the left plot, we set $\frac{z^0(i^*)}{\|z^0\|_1} = \frac{1}{s}$ – this corresponds to an "average size"-column of Z_0. In the right, we instead

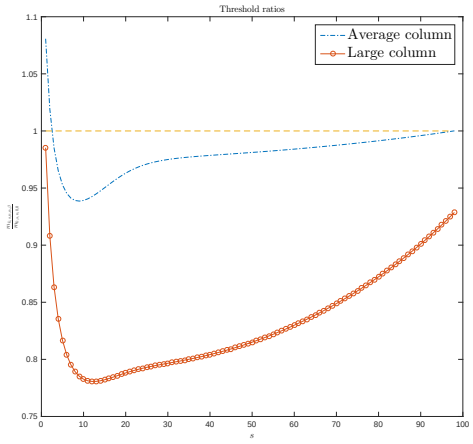

Figure 4: The relation between $m_{k,s,n,\alpha,\beta}$ and $m_{k,s,n,0,0}$ for $k = 10$, $n = 100$, different values of s and two types of columns.

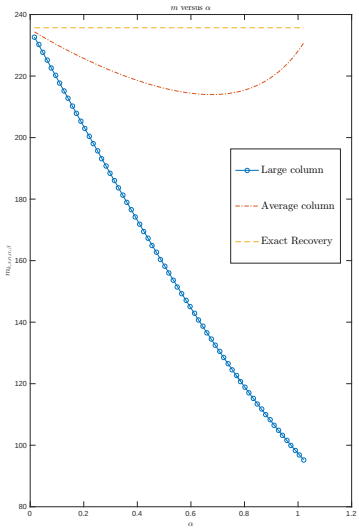

Figure 5: $m_{s,n,k,\alpha,\beta}$ for $k = s = 10$, $n = 100$ depending on α for two types of columns.

fix $\frac{z^0(i^*)}{\|z^0\|_1} = .9$ – this corresponds to a large column. With the exception of $s = 1$ (for which $1/s > .9$), $m_{k,s,n,\alpha,\beta}$ is larger for the average size column case than for the large column case, as is intuitively clear - it should be easier to recover (the more significant) large columns. We see that the isolines in both plots resemble hyperbolas $s \cdot k = \text{const.}$, which indicates that $m_{k,s,n,\alpha,\beta}$ scales with $s \cdot k$ rather than $s + k$ - that indicates soft recovery needs asymptotically as many measurements as exact recovery. This is also not surprising - after all, the results obtained in this section do not use any low-rank assumptions at all.

In practice however, also non-asymptotic reductions are relevant. This aspect is dealt with in Figure 4, where we fix $k = 10$, $n = 100$, $\alpha = \frac{\pi}{10}$ and plot the quotient $\frac{m_{k,s,n,\alpha,\beta}}{m_{k,s,n,0,0}}$ - for an average size and large column, respectively. Note that $m_{k,s,n,0,0}$ exactly corresponds to the threshold amount of measurements needed for exact recovery. Also note that since $m_{k,s,n,\alpha,\beta}$ only is an upper bound on the amount of measurements needed, it is not a contradiction that it is larger than $m_{k,s,n,0,0}$ for small values of s. We see that in particular in the case of a large column, we need considerably less measurements to ensure soft recovery than to ensure exact recovery.

Finally, in Figure 5, we fix $k = s = 10$ and $n = 100$ and plot $m_{k,s,n,\alpha,\beta}$ depending on the size of α for an average size and large column. We see that at least for small α, our result provides a smaller upper bound on the measurements needed for soft recovery than for exact recovery. Note that $m_{k,s,n,\alpha,\beta}$ is growing for large α (for really large α, it even surpasses $m_{k,s,n,0,0}$. This is again not a contradiction - we have only provided an upper bound, and we already addressed that some of the estimation we make become worse as α (and therefore also β) grows.

3.2 Energy Concentration

The condition for soft recovery described above is not quite enough to secure that not only the directions, but also the magnitudes, of the columns in the minimizer $\widehat{Z}$ are close to the ones in Z_0, but almost. The following proposition holds.

Proposition 3.6. *Let $Z_0 = z_0.H_0$ be supported on the set S and $\hat{z}.\widehat{H}$ be the minimizer of the program $\mathcal{P}_{1,2}$ with $b = \mathcal{A}(z_0.H_0)$. Assume that for each $i \in S$, $\omega(h_i^0, \hat{h}_i) \leq \alpha$, that there for some i^* exists a vector $p \in \mathbb{R}^m$ with the*

properties (3)-(6)*, and additionally*

$$\|A_i^* p\|_2 \leq \gamma, \quad i \notin S \tag{10}$$

for some $\gamma < 1$*. Then the vector* $\hat{z}$ *obeys*

$$\|\hat{z}_{S^c}\|_1 \leq \frac{1-\cos\alpha}{\cos\alpha(1-\gamma)}\hat{z}_{i^*} \tag{11}$$

and

$$\|\hat{z}_S - z_S^0\|_2 \leq \sin\alpha \frac{\max_{i\in S}\|A_i\|}{\sigma_{\mathbb{R}^S}(A_{H_0})}\|z_0\|_1 + \frac{\max_{i\in[n]}\|A_i\|}{\sigma_{\mathbb{R}^S}(A_{H_0})}\|\hat{z}_{S^c}\|_1. \tag{12}$$

$\sigma_{\mathbb{R}^S}(A_{H_0})$ *denotes the to* $\mathbb{R}^S$ *restricted singular value of* A_{H_0}*, i.e.*

$$\sigma_{\mathbb{R}^S}(A_{H_0}) = \min_{\|x\|_2=1, \operatorname{supp} x \subseteq S} \|A_{H_0}(x)\|_2,$$

and $\|\cdot\|$ *denotes the operator norm.*

Proof. Let us begin by proving (11). Since $\hat{z}.\widehat{H}$ obeys the linear constraint, $z_0.H_0 - \hat{z}.\widehat{Z} \in \ker\mathcal{A} = \operatorname{ran}\mathcal{A}^\perp$. Consequently, $\left\langle \mathcal{A}^* p, z_0.H_0 - \hat{z}.\widehat{H}\right\rangle = 0$, i.e.

$$\|z_0\|_1 \leq \sum_{i\in S}\left\langle A_i^* p, z_i^0 h_i^0\right\rangle = \sum_{i\in[n]}\left\langle A_i^* p, \hat{z}_i\hat{h}_i\right\rangle \leq \tfrac{1}{\cos\alpha}\hat{z}_{i^*} + \sum_{i\in S\setminus\{i^*\}}\hat{z}_i + \gamma\sum_{i\notin S}\hat{z}_i = \left(\tfrac{1}{\cos\alpha}-1\right)\hat{z}_{i^*} + \|\hat{z}_S\|_1 + \gamma\|\hat{z}_{S^c}\|_1$$

We used (3), (6) and (10). At the same time, due to the optimality of $\hat{z}.\widehat{H}$, there must be $\|\hat{z}_S\|_1 + \|\hat{z}_{S^c}\|_1 \leq \|z_0\|_1$. These two inequalities imply

$$\cos\alpha(1-\gamma)\|\hat{z}_{S^c}\|_1 \leq (1-\cos\alpha)\hat{z}_{i^*}$$

which implies the first half of (11).

To prove the second estimate, we again utilize that $\mathcal{A}(z_0.H_0 - \hat{z}.\widehat{H}) = 0$. That equation can namely be written

$$\begin{aligned}
\sum_{i\in S} A_i^* h_i^0(z_i^0 - \hat{z}_i) + \sum_{i\in S} A_i^*(h_i^0 - \hat{h}_i)\hat{z}_i - \sum_{i\notin S}\hat{z}_i A_i^*\hat{h}_i = 0 \Rightarrow \\
\|A_{H_0}(z_0 - \hat{z})\|_2 \leq& \sum_{i\in S}\hat{z}_i\|A_i^*(h_i^0 - \hat{h}_i)\| + \sum_{i\notin S}\hat{z}_i\|A_i^*\hat{h}_i\| \\
\leq& \|\hat{z}_S\|_1 \max_{i\in S}\|A_i\|\sin(\alpha) + \|\hat{z}_{S^c}\|_1\max_{i\notin S}\|A_i\| \\
\leq& \sin(\alpha)\max_{i\in S}\|A_i\|\|z_0\|_1 + (1-\sin\alpha)\max_{i\in[n]}\|A_i\|\|\hat{z}_{S^c}\|_1
\end{aligned}$$

We utilized that $\|h_i^0 - \hat{h}_i\|_2 \leq \sin\alpha$ for $i \in S$ and $\|\hat{z}_S\| = \|\hat{z}\|_1 - \|\hat{z}_{S^c}\|_1 \leq \|z_0\|_1 - \|\hat{z}_{S^c}\|$. To finish the proof, we simply have to note that $\|A_{H_0}(z_0 - \hat{z}_S)\|_2 \geq \sigma_{\mathbb{R}^S}(A_{H_0})\|(z_0 - \hat{z})_S\|_2$. □

Remark 3.7. *It is known that with high probability,* $\max_{i\in[n]}\|A_i\| \lesssim \sqrt{m} + \sqrt{k}$ *and* $\sigma_{\mathbb{R}^S}(A_{H_0}) \gtrsim \sqrt{m} - \sqrt{s}$ *[8, Th. 9.6]. Hence, if we choose* $m \geq C^2\max(s,k)$*, we can guarantee that*

$$\frac{\max_{i\in[n]}\|A_i\|}{\sigma_{\mathbb{R}^S}(A_{H_0})} \leq \frac{C+1}{C-1}$$

with high probability.

3.3 Taking Low Rank Into Account.

Until know, all theory has been valid for arbitrary column-sparse matrices. In the following, we will prove that under the assumption that Z_0 has been softly recovered by the program $\mathcal{P}_{1,2}$ and Z_0 is a certain type of an r-rank matrix (we will specify the exact requirements on Z_0 later), the space spanned by the leading r left singular vectors of the recovered matrix will be close to the true range of Z_0. The argumentation will rely heavily on a classical result from pertubation theory; the so-called $\sin\theta$-theorem. We will start this section by recalling that theorem. We follow the original paper [16].

Let two matrices A and $B = A + T \in \mathbb{R}^{m,n}$ be given. Write the singular value decomposition of A in the following manner:

$$A = A_1 + A_2 = U_1(A)\Sigma_1(A)V_1(A)^* + U_2(A)\Sigma_2(A)V_2(A)^*$$

where the left and right singular vectors of A are given by the columns of the matrices $[U_1(A), U_2(A)]$ and $[V_1(A), V_2(A)]$, respectively, and the union of the diagonals of Σ_1 and Σ_2 is the set of singular values of A. The dimensions of the matrices are as follows:

$$\begin{aligned} U_1(A) \in \mathbb{R}^{m,r}, \quad \Sigma_1(A) \in \mathbb{R}^{r,r}, \quad V_1(A) \in \mathbb{R}^{n,r} \\ U_2(A) \in \mathbb{R}^{m,m-r}, \quad \Sigma_2(A) \in \mathbb{R}^{m-r,m-r}, \quad V_2(A) \in \mathbb{R}^{n,m-r} \end{aligned}$$

B_1, B_2, $U_1(B)$, $\Sigma_1(B)$ and so on are defined in the same manner. The $\sin\theta$-theorem is a statement about the principal angle between the subspaces spanned by $U_1(A)$ and $U_1(B)$ (and $V_1(A)$ and $V_1(B)$). The principal angle $\angle(E, F) \in [0, \frac{\pi}{2}]$ between two subspaces E and F are thereby defined through

$$\sin\angle(E, F) = \|\Pi_{E^\perp}\Pi_F\|.$$

Theorem 3.8 (The $\sin\theta$-theorem)**.** *[16] Assume the matrices A, B and T are given as above. Define the parameter μ through*

$$\mu = \max(\|\Pi_{\operatorname{ran} B_1} T\|, \|T\Pi_{\ker B_1^\perp}\|).$$

Further assume that there exists a $\tau \geq 0$ and a $\delta > 0$ such that

$$\sigma_{\min}(A_1) \geq \tau + \delta \ \wedge \ \sigma_{\max}(B_2) \leq \tau.$$

Then there holds

$$\max(\sin\angle(\operatorname{ran} A_1, \operatorname{ran} B_1), \sin\angle(\ker A_1^\perp, \ker B_1^\perp)) \leq \frac{\mu}{\delta}.$$

With the above theory, we can prove the following result. We use the notation $\langle A\rangle := \operatorname{ran} A$.

Proposition 3.9. *Let the matrix $Z_0 = z_0.H_0$ be supported on the set S and have at most rank r. Assume furthermore that S can be partitioned into L disjoint subsets S_ℓ, $\ell = 1, \ldots, L$ so that*

$$h_i^0 = \eta_\ell^0 \mod \pm 1, \quad i \in S_\ell$$

for some unit-norm frame $(\eta_\ell)_{\ell\in[L]}$ of $\langle H_0\rangle$ with lower frame bound Λ.

If then $\widehat{Z} = \hat{z}.\widehat{H}$ has the properties that $\omega(h_i^0, \hat{h}_i) \leq \alpha < \frac{\pi}{2}$, we have for the space $\langle\widehat{H}\rangle_r$ spanned by the r first singular vectors of $\widehat{H}$

$$\sin\angle(\langle\widehat{H}\rangle_r, \langle H_0\rangle) \leq \frac{\max(\sin\alpha\|\hat{z}_S\|_2, \|\hat{z}_{S^c}\|_2)}{\sqrt{\Lambda\min_\ell\|\hat{z}_{S_\ell}\|_2^2 - \sin\alpha\|\hat{z}_S\|_2^2}}$$

Proof. Let $\epsilon > 0$ be arbitrary and define the matrices B and T through

$$B(i) = \begin{cases} \max(\hat{z}_i, \epsilon) \left\langle \hat{h}_i, h_i^0 \right\rangle h_i^0 & i \in S \\ 0 & i \notin S \end{cases} \quad T(i) = \begin{cases} \hat{z}_i \hat{h}_i - \max(\hat{z}_i, \epsilon) \left\langle \hat{h}_i, h_i^0 \right\rangle h_i^0 & i \in S \\ \hat{z}_i \hat{h}_i & i \notin S \end{cases}$$

Adopting the notation from above, we further set $A = \hat{z}.\widehat{H}$. Then $A = B + T$ and, continuing to use the notation from above, $\langle A_1 \rangle = \langle \widehat{H} \rangle_r$, $\langle B_1 \rangle = \langle H_0 \rangle$. (The latter is due to $\max(\hat{z}_i, \epsilon) \langle \hat{h}_i, h_i^0 \rangle \neq 0$ for all i, which is a consequence of $\alpha < \frac{\pi}{2}$ and $\epsilon > 0$.) Hence

$$\sin \angle(\langle H_0 \rangle, \langle \widehat{H} \rangle_r,) = \sin \angle(\langle A_1 \rangle, \langle B_1 \rangle).$$

We now want to apply the $\sin\theta$-theorem. Towards this end, we have to estimate the parameters μ, τ and δ. τ can be chosen equal to zero – the range of B is contained in $\langle H_0 \rangle$, and hence is at most r-dimensional. Consequently, B has at most r non-zero singular values, and $B_2 = 0$. To estimate μ, we begin by noting that $\ker B_1^\perp = \ker B^\perp \subseteq \mathbb{R}^S$ (since $\mathbb{R}^{S^c} \subseteq \ker B$). This implies

$$\|T\Pi_{\ker B_1^\perp}\| \leq \|T\Pi_{\mathbb{R}^S}\| \leq \sup_{\|x\|_2 \leq 1} \sum_{i \in S} |x_i| (\epsilon |\langle \hat{h}_i, h_i^0 \rangle| + \hat{z}_i \|\Pi_{\langle H_0 \rangle^\perp} \hat{h}_i\|_2) \leq \sqrt{s}\epsilon + \|\hat{z}_S\|_2 \sin\alpha.$$

In the last step, we used the Cauchy Schwarz inequality and

$$\|\Pi_{\langle H_0 \rangle^\perp} \hat{h}_i\|_2 = \min_{v \in \langle H_0 \rangle} \|\hat{h}_i - v\|_2 \leq \|\hat{h}_i - \langle \hat{h}_i, h_i^0 \rangle h_i^0\|_2 = \sqrt{1 - \cos^2(\omega(\hat{h}_i, h_i^0))} \leq \sin\alpha$$

due to $\omega(\hat{h}_i, h_i^0) \leq \alpha$ for $i \in S$. We also have, since $\langle B_1 \rangle = \langle H_0 \rangle$

$$\|\Pi_{\operatorname{ran} B_1} T\| = \|\hat{z}_{S^c}.(\Pi_{\langle H_0 \rangle} \widehat{H})\| \leq \|\hat{z}_{S^c}\|_2.$$

Here we used that the vectors in $\widehat{H}$ have unit norm and an argument similar as above. Summarizing,

$$\mu \leq \max(\sqrt{s}\epsilon + \sin\alpha \|\hat{z}_S\|_2, \|\hat{z}_{S^c}\|_2)$$

It remains to estimate the gap δ. Since we already have noted that τ can be chosen equal to zero we mearly need to estimate $\sigma_{\min}(A_1) = \sigma_r(A)$ from below. To do this, we use two well known facts: firstly the equality $\sigma_r(A) = \sigma_r(A^*)$, and secondly the so-called max-min-principle (or *Courant-Fischer Theorem*):

$$\sigma_r(A^*) = \max_{\dim V \leq r} \min_{\substack{p \in V \\ \|p\|_2 = 1}} \|A^* p\|.$$

Hence, since $\dim \langle H_0 \rangle \leq r$

$$\sigma_r(A^*) \geq \min_{\substack{p \in \langle H_0 \rangle \\ \|p\|_2 = 1}} \|(\hat{z}.\widehat{H})^* p\|_2.$$

The adjoint of $\hat{z}.\widehat{H}$ is given by $((\hat{z}.\widehat{H})p)_i = \hat{z}_i \langle \hat{h}_i, p \rangle$. Hence

$$\|(\hat{z}.\widehat{H})^* p\|_2^2 = \sum_{i \in [n]} \hat{z}_i^2 |\langle \hat{h}_i, p \rangle|^2 \geq \sum_{i \in S} \hat{z}_i^2 |\langle \hat{h}_i, p \rangle|^2$$

Now we notice that if $\|p\|_2 = 1$, the following is true due to the triangle inequality of ω and $\omega(\hat{h}_i, h_i^0) \leq \alpha$

$$|\langle \hat{h}_i, p \rangle|^2 = \cos^2(\omega(\hat{h}_i, p)) \geq \cos^2(\min(\omega(h_i^0, p) + \alpha, \tfrac{\pi}{2})) \geq \cos^2(\omega(h_i^0, p)) - \sin\alpha = |\langle h_i^0, p \rangle|^2 - \sin\alpha,$$

where the proof of the last inequality is postponed to the Appendix (more specifically Lemma A.3). Therefore, we may estimate

$$\sum_{i\in S} \hat{z}_i^2 |\langle \hat{h}_i, p\rangle|^2 \geq \sum_{i\in S} \hat{z}_i^2 |\langle h_i^0, p\rangle|^2 - \sin(\alpha)\|\hat{z}_S\|_2^2$$

Now we utilize the structure of the set $(\hat{h}_i)_{i\in S}$ to estimate

$$\sum_{i\in S} \hat{z}_i^2 |\langle h_i^0, p\rangle|^2 = \sum_{\ell\in[L]} |\langle \eta_\ell^0, p\rangle|^2 \sum_{i\in S_\ell} \hat{z}_i^2 \geq \min_\ell \|z_{S_\ell}\|_2^2 \sum_{\ell\in[L]} |\langle \eta_\ell^0, p\rangle|^2 \geq \Lambda \min_\ell \|z_{S_\ell}\|_2^2$$

for all $p \in \langle H_0\rangle$ with $\|p\|_2$. All in all,

$$\delta^2 \geq \Lambda \min_\ell \|z_{S_\ell}\|_2^2 - \sin(\alpha)\|\hat{z}_S\|_2^2,$$

which together with the $\sin\theta$-theorem proves

$$\sin\angle(\langle \widehat{H}\rangle_r, \langle H_0\rangle) \leq \frac{\max(\epsilon\sqrt{s} + \sin\alpha\|\hat{z}_S\|_2, \|\hat{z}_{S^c}\|_2)}{\sqrt{\Lambda \min_\ell \|\hat{z}_{S_\ell}\|_2^2 - \sin\alpha\|\hat{z}_S\|_2^2}}.$$

Letting $\epsilon \to 0$ yields the claim. □

Remark 3.10. *In the case that $r = 1$, all columns h_i^0, $i \in S$ must be equal modulo sign, which implies that we may choose $L = 1$, $\eta_1^0 = h_1^0$, $S_1 = S$ and consequently $\Lambda = 1$. We then arrive at the cleaner estimate*

$$\sin\angle(\langle \widehat{H}\rangle_r, \langle H_0\rangle) \leq \frac{\max(\sin\alpha\|\hat{z}_S\|_2, \|\hat{z}_{S^c}\|_2)}{\|\hat{z}_S\|_2\sqrt{1-\sin\alpha}}.$$

4 Heuristic Proposals How One Could Utilize Soft Recovery

In this section, we will present two ideas how one could use the phenomenon of soft recovery to design a recovery algorithm for matrices which are both column-sparse and have low rank. These ideas are of highly heuristic nature, whence we mostly study them experimentally.

4.1 Thresholding

The above analysis indicates that the set of indices where $\hat{z}_i$ is large probably coincides relatively well with the support of z_0, since $\|\hat{z}_{S^c}\|_1$ is small. Hence, if we choose a set $\widehat{S}$ for which $\|\widehat{Z}_{\widehat{S}}\|_{1,2} \geq \|\widehat{Z}\|_{1,2} * \tau$ for $\tau \approx 1$ and $\widehat{S} \geq s$, $\hat{S}$ is most probably a small set which still contains the support of the ground truth signal Z_0. After having identified the support, the remaining task is then to recover a low-rank matrix of dimension $k \times |\widehat{S}|$, which we choose to do with nuclear norm-minimization minimization. This will be successful with high probability already when $m \gtrsim r(k + |\widehat{S}|)\log(k|\widehat{S}|)$ [14].

The procedure outlined above is summarized in Algorithm 1.

4.1.1 Numerical Experiments

To test the performance of the $NAST$-algorithm, we perform a small numerical experiment. For $m = 100, 105, \ldots, 300$, we generate random 10-column sparse matrices in $\mathbb{R}^{10,100}$ of the form $\sum_{i=1}^r h_i z_i^*$ for $r = 1, 2$ and 5 (in all experiments $s = k = 10$, $n = 100$.) The vectors h_i are independent and uniformly distributed on the sphere where as z has a uniformly drawn random support and normally distributed and independent non-zero entries. We measure these matrices with a randomly drawn Gaussian measurement matrix, and test if $NAST$ is able to recover the ground truth signal. The minimization problems were solved with help of the $MATLAB$-package `cvx` [12]. A success is declared if the relative error in Frobenius norm is less than 0.1%. This experiment was repeated a

Algorithm 1: (NAST) Nuclear norm After Soft recovery Thresholding.

Data: A linear map $\mathcal{A} : \mathbb{R}^{k,n} \to \mathbb{R}^q$, a sparsity s, a parameter $\tau \in (0, 1)$ and a vector b.
Result: An estimate Z_* of a sparse solution of $\mathcal{A}(Z) = b$, where $Z_* = z_*.H_*$ with z_* sparse and H_* low rank.

1 Solve the $\mathcal{P}_{1,2}$ for a column-sparse matrix $\widehat{Z} = \hat{z}.\widehat{H}$.
2 $\widehat{S} \leftarrow$ smallest set S with at least s indices so that $\|\widehat{Z}_S\|_{1,2} \geq \tau \|\widehat{Z}\|_{1,2}$.
3 Solve the nuclear minimization problem

$$\min_{\operatorname{supp} Z \subseteq \widehat{S}} \|Z\|_* \text{ subject to } \mathcal{A}(Z) = b$$

and output the solution Z_*.

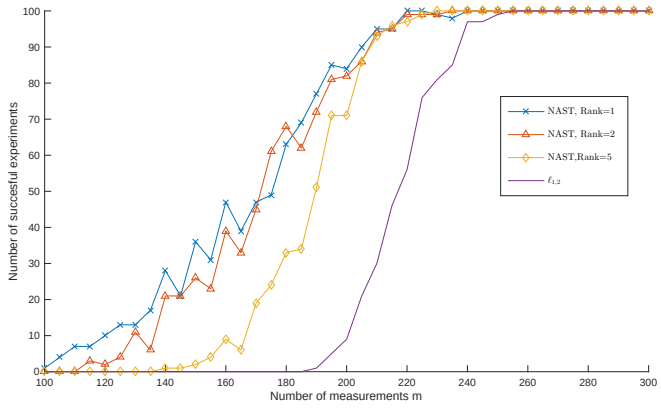

Figure 6: Results of numerical performance tests of $NAST$.

hundered times, and the results are depicted in Figure 6. We can clearly see that the $NAST$-procedure outperforms $\ell_{1,2}$. We also see that the performance is better for low ranks, as expected. We performed a small control experiment for $m = 300$ testing the performance of nuclear norm minimization - out of a hundered trials, there was not a single success.

4.2 "Column Streamlining"

Another idea, which is a bit more involved, we want to present is the following: If Z_0 is of rank r, the columns containing most of the energy of $\widehat{Z}$ all almost lie in an r-dimensional subspace, not far from $\operatorname{ran} Z_0$. Hence, the best rank r-approximation V of $\operatorname{ran} \widehat{Z}$ should be close to $\operatorname{ran} Z_0$.

This information can be used to modify the $\ell_{1,2}$-norm to more greatly penalize components of the norms of Z not lying in the space V. We propose the following way of doing this

$$\min \|Z\| + \|\Pi_{V^\perp} Z\|_{1,2} \text{ subject to } \mathcal{A}(Z) = b.$$

The procedure of alternately calculating matrices Z and subspaces V according to the procedure described above is then repeated until some stopping criterion is met. Some heuristic proposals for this criterion are for instance that the difference between iterate Z_k and Z_{k+1} drops below some threshold ϵ, or the same for the $(r + 1)$:st singular value of Z_k (indicating that we have found a low-rank solution), or also the $(s + 1)$:st largest column (indicating that we have found a sparse solution). Alternatively, one could break the iteration already when, say, $2s$ columns are larger than some threshold, form a set $\widehat{S}$ of the corresponding indices, and solve a low-rank

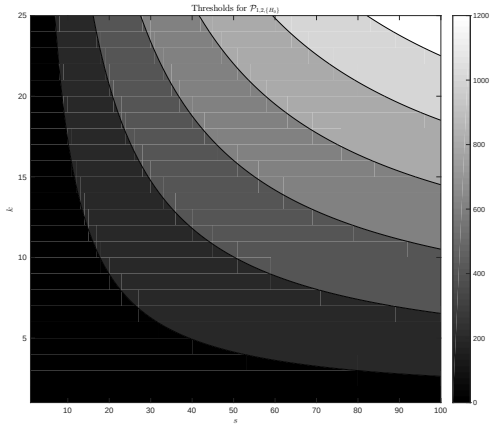

Figure 7: Plot of the function (13) for $n = 100$ and $r = 1$.

problem with that constraint, as above. The main algorithm is summarized in Algorithm 2. We have chose to call it the *Column Streamlining* algorithm, since it, when successful, forces the non-zero columns of the iterates to all align in a common, low-dimensional subspace.

Algorithm 2: Column Streamlining

Data: A linear map $\mathcal{A} : \mathbb{R}^{k,n} \to \mathbb{R}^m$, a rank r and a vector b.

Result: An estimate Y_* of a sparse solution of $\mathcal{A}(Y) = b$, where $Y_* = z.H$ with z sparse and H of rank less than r.

Initalize Y_0 as the solution of $\mathcal{P}_{1,2}$ for $\mathcal{A}$ and b, and V as the the best r-dimensional approximation of $\operatorname{ran} Y_0$, i.e

$$V = \operatorname{span}(U_1, \dots U_r),$$

where $U\Sigma V^*$ is the SVD of Y_0.

repeat

1 $Y_q \leftarrow \operatorname{argmin} \|Y\|_{1,2} + \|\Pi_{V_{q-1}^\perp} Y\|_{1,2}$ subject to $\mathcal{A}(Y) = b$.

2 $V_q \leftarrow$ The best r-dimensional approximation of $\operatorname{ran} Y_k$.

until *A stopping criterion is met*;

We cannot report much success about theoretical guarantees about the performance of the Column Streamlining algorithm. We however believe that it again will not perform well using the optimal amount $r(s + k)$ measurement, but instead again get stuck at the $(s \cdot k)$–bottleneck. The argument goes as follows: Let us assume that the iterates $Y_q \to Z_0$. Then it is not hard to prove that Z_0 has to solve the minimization problem

$$\operatorname{argmin} \|Z\|_{1,2} + \|\Pi_{\langle H_0\rangle^\perp} Z\|_{1,2} \text{ subject to } \mathcal{A}(Z) = b. \qquad (\mathcal{P}_{1,2,\langle H_0\rangle})$$

By following the exact same route as in Section 3.1 one may now easily prove that an upper bound of the amount of Gaussian measurements needed for the success of the program $\mathcal{P}_{1,2,\langle H_0\rangle}$ is given by

$$\mu_{k,s,n,r} = \inf_{\tau>0} s(r + \tau^2 + \mathbb{E}\left(\operatorname{pos}(\|g^{k-r}\|_2 - \tau)^2\right)) + (n - s)\mathbb{E}\left(\operatorname{pos}(\|g^k\|_2 - \tau)^2\right). \qquad (13)$$

(The only hard part of the argument, which is to calculate the subdifferential of $\|\cdot\|_{1,2} + \|\Pi_{\langle H_0\rangle^\perp} \cdot\|_{1,2}$, is presented in Section A.1 of the Appendix.) By plotting this function in Figure 7 for $n = 100$, $r = 1$ for varying values of s

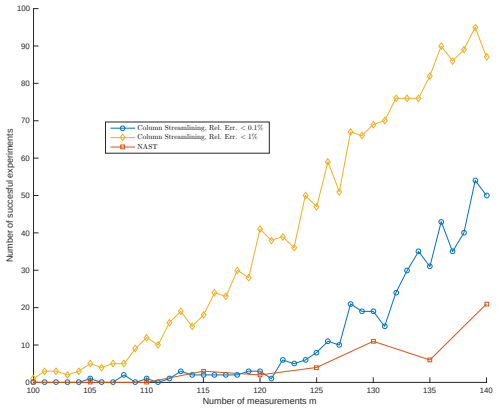
Column Streamlining, Rel. Err. < 0.1%
Column Streamlining, Rel. Err. < 1%
NAST
Number of succesful experiments
Number of measurements m

Figure 8: Numerical results of the Column Streamlining Algorithm.

and k, we see that it scales as $s \cdot k$ rather than as $s + k$. It should however be noted that in absolute numbers, we need considerably less measurements to secure success of $\mathcal{P}_{1,2,\langle H_0 \rangle}$ than for $\mathcal{P}_{1,2}$ - as an example, $\mu_{10,10,100,1} \approx 130$ whereas $m_{10,10,100,0} \approx 236$

4.2.1 Numerical Experiments

We chose to test the numerical performance of the Column Streamlining algorithm as follows: Matrices were generated as above (i.e. $k = s = 10$, $n = 100$) for $r = 2$, for each $m = 100, 101, \ldots, 140$. Then we let the Column Streamlining algorithm perform 10 iterations, and subsequently recorded the relative difference in Frobenius norm between the final iterate Y_* and the ground truth signal Z_0. This was repeated a hundered times for each value of m. We again used `cvx` to solve the minimization programs. The number of experiments in which the final relative error was smaller than 0.1% and 1%, respectively, are depicted and compared with the results of the $NAST$-experiments in Figure 8. We see that the Column Streamlining outperforms the $NAST$-approach. The big difference between the number of final iterates with a relative error smaller than 0.1% and 1% suggest that better results if we let the algorithm perform more iterations.

Acknowledgement

The author acknowledges support from Deutsche Forschungsgemeinschaft (DFG) Grant KU 1446/18 - 1 and Berlin Mathematical School. He also likes to thank Gitta Kutyniok, Saeid Haghighatshoar and Peter Jung for interesting discussions on this and other topics.

References

[1] A. Ahmed, B. Recht, and J. Romberg. Blind deconvolution using convex programming. *IEEE Trans. Inf. Theory*, 60(3):1711–1732, 2014.

[2] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: phase transitions in convex programs with random data. *Information and Inference*, 3:224 – 294, 2014.

[3] S. Boyd and L. Vanderberghe. *Convex Optimization*. Cambridge University Press, 7th edition, 2009.

[4] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inf. Theory*, 51:4203–4215, 2005.

[5] D. Donoho and J. Tanner. Sparse nonnegative solution of underdetermined linear equations by linear programming. *Proc. Natl. Acad. Sci. USA*, 102:9446–9451, 2005.

[6] Y. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE Trans. Inf. Theory*, 55(11):5302–5316, Nov 2009.

[7] A. Flinth and G. Kutyniok. *PROMP*: A Sparse Recovery Approach to Lattice-Valued Signals. 2016. Preprint.

[8] S. Foucart and H. Rauhut. *A mathematical introduction to Compressed Sensing*. Birkhäuser, 2013.

[9] Y. Gordon. On Milman's inequality and random subspaces which escape through a mesh in $\mathbb{R}^n$. In J. Lindenstrauss and V. D. Milman, editors, *Geometric Aspects of Functional Analysis*, volume 1317 of *Lecture Notes in Mathematics*, pages 84–106. Springer Berlin Heidelberg, 1988.

[10] J. P. Haldar and D. Hernando. Rank-constrained solutions to linear matrix equations using powerfactorization. *IEEE Signal Processing Lett.*, 16(7):584–587, 2009.

[11] K. Lee, Y. Wu, and Y. Bresler. Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization. *arXiv preprint arXiv:1312.0525*, 2013.

[12] M. MGrant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, Mar. 2014.

[13] S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Trans. Inf. Theory*, 61(5):2886–2908, 2015.

[14] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.

[15] M. Stojnic, F. Parvaresh, and B. Hassibi. On the reconstruction of block-sparse signals with an optimal number of measurements. *IEEE Trans. Signal Process.*, 57(8):3075–3085, 2009.

[16] P.-Å. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.

[17] P. Wojtaszczyk. Stability and instance optimality for gaussian measurements in compressed sensing. *Found. Comput. Math.*, 10(1):1–13, 2010.

A Appendix

We here provide minor proofs not contained in the main body of the article. We begin with the argument advertised in the introduction about the impossibility of soft recovery via ℓ_1-minimization.

Proposition A.1. *Let $A \in \mathbb{R}^{m,n}$ and $z_0 \in \mathbb{R}^n$ be arbitrary. If ℓ_1-minimization problem $\mathcal{P}_1$ with $b = Az_0$ has a unique solution z_*, one of the following alternatives hold*

- $z_0 = z_*$
- $z_*(i) = 0$ *for at least one* $i \in \operatorname{supp} z_0$.

Proof. Denote $\operatorname{supp} z^* = S^*$. Towards a contradiction, assume that $z_0 \neq z_*$ and that the support S of z_0 is contained in S_*. By [7, Lem. 5.1] (which little more than a specification of [5, Thm. 1]), $\frac{z_*}{\|z_*\|_1}$ then necessarily lies in the relative interior of an $(|S^*|-1)$-face F of the crosspolytope

$$C_1 = \{x \in \mathbb{R}^n \mid \|x\|_1 \leq 1\}$$

which is mapped to an $(|S^*|-1)$-face of AC_1. In particular, the columns (a_i) of A corresponding to $i \in S^*$ are linearly independent. But due to the linear constraint, we have

$$b = \sum_{i \in S} z_0(i) a_i = \sum_{i \in S^*} z_*(i) a_i.$$

This equation, together with $S \subseteq S_*$ and the linear independence, readily implies $z_0 = z_*$. □

Next, we calculate the dual problem of $\mathcal{P}_1^+$, as was needed in the proof of Lemma 2.2. Note that this calculation is relatively standard, but we include it for completeness.

Lemma A.2. *The dual problem of*

$$\min \|z\|_1 = \sum_{i \in [n]} z_i \text{ subject to } Ax = b \ , \ x \geq 0$$

is equivalent to

$$\max \langle b, p \rangle \text{ subject to } A^* p \leq \mathbb{1},$$

where $\mathbb{1}$ *is the vector consisting only of entries equal to* 1

Proof. Let us begin by writing the Lagrange dual

$$\mathcal{L}(z, p, \mu) = \|z\|_1 + \langle p, b - Az \rangle - \langle \mu, z \rangle = \langle b, p \rangle + \|z\|_1 - \langle z, A^* z + \mu \rangle$$

We see that $\inf_z \mathcal{L}(z, p, \mu)$ is finite exactly when $\|A^* p + \mu\|_\infty \leq 1$, in which case it is equal to $\langle b, p \rangle$. The dual problem is hence

$$\max_{p, \mu \geq 0} \langle b, p \rangle \text{ subject to } \|A^* p + \mu\|_\infty \leq 1.$$

Due to the fact that there exists a $\mu \geq 0$ with $\|A^* p + \mu\| \leq 1$ if and only if $A^* p \leq \mathbb{1}$, we see the stated equivalence. □

Finally, we prove the trigonometric inequality used in the proof of Proposition 3.9.

Lemma A.3. *For* $x \in [0, \pi]$ *and* $\alpha \in [0, \frac{\pi}{2}]$*, we have*

$$\cos^2 \left(\min \left(x + \alpha, \tfrac{\pi}{2} \right) \right) \geq \cos^2(x) - \sin \alpha$$

Proof. We treat the cases $x + \alpha \leq \frac{\pi}{2}$ and $x + \alpha \geq \frac{\pi}{2}$ separately. In the first case, we have

$$\begin{aligned}
\cos^2 \left(\min \left(x + \alpha, \tfrac{\pi}{2} \right) \right) &= \cos^2(x + \alpha) = (\cos(x)\cos(\alpha) - \sin(x)\sin(\alpha))^2 \\
&= \cos^2(x)\cos^2(\alpha) - 2\cos(x)\sin(x)\cos(\alpha)\sin(\alpha) + \sin^2(x)\sin^2(\alpha) \\
&= \cos^2(x) - \sin(\alpha)\left((\cos^2(x) - \sin^2(x))\sin(\alpha) + 2\cos(x)\sin(x)\cos(\alpha) \right) \\
&= \cos^2(x) - \sin(\alpha)\left(\cos(2x)\sin(\alpha) + \sin(2x)\cos(\alpha) \right) \\
&= \cos^2(x) - \sin(\alpha)\sin(2x + \alpha) \geq \cos^2(x) - \sin(\alpha).
\end{aligned}$$

In the second case, we may argue as follows: α and $\frac{\pi}{2}$ both lie in the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$, in which sin is increasing. The inequality $\alpha \geq \frac{\pi}{2} - x$ therefore implies $\sin(\alpha) \geq \sin\left(\frac{\pi}{2} - x\right) = \cos(x)$. Hence

$$\cos^2(x) - \sin(\alpha) \leq \cos^2(x) - \cos(x) \leq 0 = \cos^2 \left(\min \left(x + \alpha, \tfrac{\pi}{2} \right) \right),$$

which proves the claim. □

A.1 The Subdifferential of the Norm Used In $\mathcal{P}_{1,2,\langle H_0\rangle}$

Here we provide the calculation of the subdifferential of $\|\cdot\|_{1,2} + \|\Pi_{\langle H_0\rangle^\perp}\cdot\|_{1,2}$ at Z_0, as was advertised in 4.2.

Proposition A.4. *Let $Z_0 = z_0.H_0$ be supported on the set S. Then the subdifferential of $\|\cdot\|_{1,2} + \|\Pi_{\langle H_0\rangle^\perp}\cdot\|_{1,2}$ at Z_0 is given by the set of matrices V for which*

$$V(i) \in h_i^0 + B_1(0) \cap \langle H_0\rangle^\perp, \quad i \in S$$
$$V(i) \in B_1(0) + B_1(0) \cap \langle H_0\rangle^\perp, \quad i \notin S.$$

Proof. We are looking for the set of matrices V for which

$$\sum_{i\in S} \|h_i^0 + U(i)\|_2 + \|\Pi_{\langle H_0\rangle^\perp} U(i)\|_2 + \sum_{i\notin S} \|U(i)\|_2 + \|\Pi_{\langle H_0\rangle^\perp} U(i)\|_2 \geq \sum_{i\in S} \|h_i^0\|_2 + \langle V(i), U(i)\rangle + \sum_{i\notin S} \langle V(i), U(i)\rangle$$

for all $U \in \mathbb{R}^{k,n}$. We used that $h_i^0 \in \langle H_0\rangle$ for all $i \in S$. Since the columns of U can be chosen independently, we can treat each index separately. In the following, we drop the indices in the calculations for notational simplification, and use lower case letters to denote columns of V and U, respectively.

Case 1: $i \in S$**.** We aim to characterize the vectors $v \in \mathbb{R}^k$ for which

$$\|h^0 + u\|_2 + \|\Pi_{\langle H_0\rangle^\perp} u\|_2 \geq \|h^0\|_2 + \langle v, u\rangle \tag{14}$$

for all vectors $u \in \mathbb{R}^k$. If $v \in h^0 + B_1(0) \cap \langle H_0\rangle^\perp$, (14) is satisfied:

$$\begin{aligned}\|h^0\|_2 + \langle v, u\rangle = \|h^0\|_2 + \langle h^0, \Pi_{\langle H_0\rangle} u\rangle + \left\langle \Pi_{\langle H_0\rangle^\perp} v, \Pi_{\langle H_0\rangle^\perp} u\right\rangle &\leq \|h^0 + \Pi_{\langle h^0\rangle} u\|_2 + \|\Pi_{\langle H_0\rangle^\perp} u\|_2 \\ &\leq \|h^0 + u\|_2 + \|\Pi_{\langle H_0\rangle^\perp} u\|_2.\end{aligned}$$

We now only need to argue that (14) implies that $v \in h^0 + B_1(0) \cap \langle H_0\rangle^\perp$. By plugging in $\pm h^0$ for u, we see that

$$\langle h^0, v\rangle = \|h^0\|_2.$$

Hence, v is of the form $h^0 + v_1 + v_2$ for $v_1 \in \langle H_0\rangle$, $v_1 \perp h^0$ and $v_2 \perp \langle H_0\rangle$. Now, by plugging in λv_1 with $\lambda > 0$ for u, we obtain

$$\|h^0\|_2 + \lambda\|v_1\|_2^2 \leq \|h^0 + \lambda v_1\|_2 = \sqrt{\|h^0\|_2^2 + \lambda^2\|v_1\|_2^2} \Rightarrow \|v_1\|^2 \leq \frac{\sqrt{\|h^0\|_2^2 + \lambda^2\|v_1\|_2^2} - \|h^0\|_2}{\lambda}, \lambda > 0$$

By letting $\lambda \to 0$, we obtain $\|v_1\|_2 = 0$. Similarly, plugging in λv_2 for u, we obtain

$$\|h^0\| + \lambda\|v_2\|_2^2 \leq \|h^0 + \lambda v_2\|_2 + \lambda\|v_2\|_2 \Rightarrow (\|v_2\|_2^2 - \|v_2\|_2) \leq \frac{\sqrt{\|h^0\|_2^2 + \lambda^2\|v_2\|_2^2} - \|h^0\|_2}{\lambda}, \lambda > 0.$$

Again by letting $\lambda \to 0$, we obtain $(\|v_2\|_2^2 - \|v_2\|_2) \leq 0$, which only is satisfied if $\|v_2\|_2 \leq 1$.

Case 2: $i \notin S$**.** Here the aim is to characterize the vectors v for which

$$\|u\|_2 + \|\Pi_{\langle H_0\rangle^\perp} u\|_2 \geq \langle v, u\rangle \tag{15}$$

is true for every $u \in \mathbb{R}^k$. It is not hard to see that (15) is satisfied for every v of the form $w_1 + w_2$ with $\|w_1\|_2 \leq 1$ and $\|w_2\|_2 \leq 1$, $w_2 \perp \langle H_0\rangle$.

To see that (15) implies that v can be written as claimed is a bit more tricky. It is clear that we can always write v in the form $\mu\eta_1 + \lambda\eta_2$ with $\eta_1 \in \langle H_0\rangle$ and $\eta_2 \perp \langle H_0\rangle$ both having unit norm. Hence, $v = (\mu\eta_1 + (\lambda-1)\eta_2) + \eta_2$ and the proof is finished as soon as we have argued that $\mu^2 + (\lambda-1)^2 \leq 1$. However, plugging in $\mu\eta_1 + (\lambda-1)\eta_2$ for u in (15) yields

$$\sqrt{\mu^2 + (\lambda-1)^2} + (\lambda - 1) \geq \mu^2 + \lambda(\lambda-1) \Rightarrow \sqrt{\mu^2 + (\lambda-1)^2} \geq \mu^2 + (\lambda-1)^2,$$

which implies $\mu^2 + (\lambda-1)^2 \leq 1$. ∎